arXiv:1611.05642v1 [cs.CR] 17 Nov 2016

Data Minimisation: a Language-Based Approach (Long Version)

Thibaud Antignac[1], David Sands[1], and Gerardo Schneider[2]

[1] Chalmers University of Technology, 412 96 Göteborg, Sweden.
{thibaud.antignac,dave}@chalmers.se
[2] Gothenburg University, 412 96 Göteborg, Sweden.
gerardo@cse.gu.se

Abstract. Data minimisation is a privacy-enhancing principle considered as one of the pillars of personal data regulations. This principle dictates that personal data collected should be no more than necessary for the specific purpose consented by the user. In this paper we study data minimisation from a programming language perspective. We assume that a given program embodies the purpose of data collection, and define a data minimiser as a pre-processor for the input which reduces the amount of information available to the program without compromising its functionality. In this context we study formal definitions of data minimisation, present different mechanisms and architectures to ensure data minimisation, and provide a procedure to synthesise a correct data minimiser for a given program.

1 Introduction

A standard interpretation of sequential programs is to interpret them as functions taking data as input, performing some computation and giving a result. The result should somehow be dependent on the input, so provided enough "good" data is given as input the program should be able to perform computations and give the expected result in conformance with its functional specification.

There are many reasons for not providing more input data than needed. Besides time and space, other non-functional requirements can benefit from minimisation of the input. This is particularly relevant when processing personal, eventually sensitive, data. According to the Article 5 of the *General Data Protection Regulation* proposal (GDPR) "Personal data must be [...] limited to what is necessary in relation to the purposes for which they are processed" [11].[3] Note that this principle exists in other regulations and is sometimes referred to as "collection limitation" when it focuses on the particular step of collecting data. This is for instance the case of the *Fair Information Practice Principles* (FIPPs) [31] in USA, and the *Guidelines on the Protection of Privacy and Transborder Flows*

[3] The *General Data Protection Regulation* (EU —2016/679) was adopted on 27 April 2016, and it will enter into application 25 May 2018. Unlike a Directive it does not require any enabling legislation to be passed by governments.

```
1 input(salary)
2 benefits := (salary < 10000)
3 output(benefits)
```

Fig. 1: Running example P_{bl} to compute a benefits level.

of Personal Data [25] proposed by the Organisation for Economic Co-operation and Development (OECD).

If we consider data minimisation from the regulatory point of view, the input data not semantically used in the program should neither be collected nor processed, given that the adversary is the data processor[4] itself, having thus the possibility to exploit the inputs.[5] As a consequence the attacker knows all the information available after the input is collected (before the program execution).

Given that *input data* = *necessary data* + *extra data*, and since the program should execute equally well without any extra data, we have that *input data* $\geq$ *necessary data*. The goal of the *data minimisation process* is thus to minimise the input data so only what is necessary is given to the program. Whenever the input data exactly matches what is necessary we may say that the minimisation is the best. Best minimisation is, however, difficult to achieve in general among other things because it is not trivial to exactly determine what is the input needed to compute each possible output.

Let us consider a simple program to exemplify the notion of data minimisation. The purpose of the program in Fig. 1 is to compute the *benefit level* of employees depending on their salary (that we assume to be between \$0 and \$100000). For the sake of simplicity, in what follows we do not assume any particular distribution over their domain for the inputs. We thus drive our analysis on worst-case assumptions. A quick analysis shows that the range of the output is $\{\texttt{false}, \texttt{true}\}$, and consequently the data processor does not need to precisely know the real salaries of the employees to determine the benefit level. In principle each employee should be able to give any number between 0 and 9999 as input if they are eligible to the benefits, and any number between 10000 and 100000 otherwise, without disclosing their real salaries.

In this paper we address the issue of data minimisation with the following contributions:

- We present different mechanisms and architectures to ensure data minimisation (Section 2),
- We provide a formal definition of the concept of *minimisers for a program* used to achieve data minimisation (Section 4) with respect to an explicit attacker model (Section 3),
- We show how we can compare minimisers, leading to the definition of *best minimiser* for a given program (Section 5), and

[4] "Data controllers" and "data processors" are legal roles used to define obligations and liabilities of the parties. We indistinctly use the term "data processor" in this paper as we are interested in designating the party that technically processes the data.

[5] In other scenarios the adversary only has access to the outputs (cf. [29]).

– We propose a (semi-)procedure to generate minimisers and prove its soundness and termination (Section 6).

In addition, we discuss a proof-of-concept implementation to compute minimisers, based on symbolic execution and the use of a SAT solver, along with two examples (Section 7). Finally we discuss related work (Section 8) before concluding and giving future research directions (Section 9).

2 The When, How, and Where of Data Minimisation

Having informally introduced what minimality and minimisation are, the natural question arising is how to achieve it. Indeed, various mechanisms can be used for data minimisation, as detailed in Section 2.1 and can be performed in various architectures, as described in Section 2.2.

2.1 Minimisation Mechanisms

The mechanisms used to test, check, verify, or enforce data minimisation appear at different stages of the software lifecycle as illustrated in Table 1.

Before deployment Three kinds of analysis to perform data minimisation or check minimality can be used before deployment: purely static, purely dynamic, or both static and dynamic. Purely static mechanisms are used to perform a verification of the minimality property or a synthesis of a minimiser, ie., a pre-processor of the input which feeds minimalised data to the program(such as the one presented in Figure 4 for the example of P_{bl} given in introduction in Figure 1). The guarantee provided by such a static analysis is extremely strong as it covers all the possible executions. To achieve this the source code of the program is needed along with its specification. However, static analysis can be expensive and may not alway be applicable. In that case one may turn to more lightweight techniques.

At the other end of the spectrum, it is possible to use purely dynamic mechanisms to perform tests of the minimality property or to refine a pre-processor to make it more minimising. In this case, the test strategy will choose some possible inputs and test whether or not some other inputs invalidate the minimality property. For example, we could test P_{bl} with some input values, such as `0`, `1000`, and `10000`; the input values `0` and `1000` lead to the same output and would thus reveal that the program has not been fully minimised. However, an absence of common ouputs does not prove the program has been minimised. Giving weaker guarantees than purely static techniques, purely dynamic techniques only need a black box executable program to run tests (along with its specification).

A mix between static and dynamic analysis can be used by performing a static analysis on a set of dynamically selected inputs (following a test strategy). For example, it would be possible to perform a formal verification only for input value `0` for P_{bl}. In this case, a result as good as the one obtained with the purely

Table 1: Coverage, input, and output of minimisation mechanisms.

Timing of mechanism	*Kind of analysis*	*Type of mechanism*	*Coverage*	*Input*	*Output*
before deployment	static	verification	all possible executions (symbolic execution)	white box program specification	minimality
		synthesis			minimiser
	static and dynamic	test	some possible executions (test strategy)		minimality
		refinement			minimiser
	dynamic	test		black box program specification	minimality
		refinement			minimiser
online	static and dynamic	verification	all executions (symbolic execution)	white box program specification current execution	minimality
		synthesis			representative
	dynamic	monitoring	some executions (previous executions)	previous executions current execution	minimality
		refinement			representative
offline	static and log	verification	all executions (symbolic execution)	white box program specification all execution logs	minimality
	log	monitoring	some executions (previous executions)	all execution logs	

static approach could be obtained for this particular input, returning the set of values $\{0, \ldots, 9999\}$ (all giving the same output). However, it does not give any guarantee for the other input values. This approach gives strong guarantees on the inputs covered, but does not cover the full domain of program inputs (contrary to purely static approaches).

After deployment, both online and offline mechanisms can be run.

Online Concerning online mechanisms, only a mix of static and dynamic analysis or a purely dynamic approach can be used. Indeed, at this stage, it becomes meaningless to perform a purely static analysis as such an analysis is needed only once (it covers all the possible executions and thus is better run before deployment as explained previously). A mixed analysis is based on the current execution. In this case, it is still possible to reach strong guarantees about minimality or to get an input representative, ie., another input value hiding the the actual input value while still preserving correctness, which ensures minimality. Indeed, while the purely static approach consisted in generating a minimiser returning a representative for all the possible inputs, only the representative corresponding to the current execution is computed here. Returning to the example of P_{bl}, a client whose salary is of `8000` will execute the minimisation mechanism before disclosing this data. Thus, a representative (for instance, `0`) will be fixed at this time to hide the actual salary to the server without changing the result of the computation. As a consequence, the computation is more efficient but still relies on a white box program along with its specification and data relative to the current execution.

If such conditions are not acceptable, it is also possible to get better efficiency at the cost of weaker guarantees by using a purely dynamic approach. This consists in testing whether or not the previous executions with similar inputs led to the same output, meaning the minimality property is not reached. Returning again to our example, let us say the salary is to be disclosed every month and there has already been a disclosure of `7000` the previous month. Instead of disclosing its current salary of `8000`, the client will reuse the previously disclosed value of `7000` as it leads to the same output value. In this case, it has to be highlighted that even if the test or the refinement is performed against all the previous executions, the minimality property cannot be guaranteed as there might be more information leaking without it being witnessed by the past behaviours.

Instead of relying on previous executions, the dynamic analysis could also rely on arbitrary values, following a test strategy, as depicted for mechanisms run before deployment. In this case, an access to the black box program is necessary instead of the previous executions. For the sake of simplicity, we do not show this case in the Table 1 as it does not appear to make an interesting case in practice.

Offline Finally, offline mechanisms rely on execution logs to audit minimality. They also allow to get guarantees as strong as previously in the mixed static

and log analysis. As for the previous kinds of analysis, this requires a white box program along with its specification. In this case, all the behaviours belonging to the execution log are verified individually in a way very similar to the online mechanism.

A weaker analysis, similar to the online dynamic one, can also be performed here, only relying on logs. This monitoring approach only consider all the executions recorded in the logs and are only able to find witnesses of non-minimality. For instance, if two executions have been logged disclosed inputs `7000` and `8000` (both leading to the same output), the auditor knows that more data than needed has been disclosed.

It could also be possible here to monitor each past execution against arbitrary values given by a monitoring strategy (similarly to test strategies). This would require an access to the black box program and would give similarly weak results. This possibility has not been added to the Table 1 as it does not seem useful in practice.

Data minimisation itself can only be performed before deployment or online. Indeed, after execution, an offline approach is only able to detect breaches in minimality. As such, it is not a proper privacy-by-design practice and should not be used alone. However, this approach is still useful as it is the more likely to be used in already existing systems (as it would be extremely ambitious to hope all systems managing personal data to be redesigned according to privacy-by-design guidelines). Another use of online and offline mechanisms is their application by an external entity, such as a data protection authority, to perform sporadic controls (with the caveat that such controls could be subject to attacks depending on prior mechanisms run on data by the server).

The way that these mechanisms are applied may depend on the system architecture. Thus they may provide different guarantees as discussed below.

2.2 Minimisation Architectures

The kind of minimality and minimisation achievable by these mechanisms depend on when and where they are performed as shown in Table 2.

For mechanisms run before deployment or online, different kinds of minimality are applicable and different privacy properties can be targeted depending on the architecture.

Kind of Minimality Depending on the architecture, different kinds of minimalities can be achieved, denoted *monolithic* minimality and *distributed* minimality. Monolithic minimality is a minimality property expressed in the context where knowledge about all input values is accessible at a centralised location where the mechanism can be applied. On the other hand, distributed minimality is used for the cases where the input values and the associated mechanism are distributed, in which case the mechanism has only the knowledge available at its location. If the mechanism is performed by a client then only monolithic minimality can be achieved. On the other hand, if the mechanism is performed by a third-party or

Table 2: Kind of minimality applicable and privacy property targeted by minimisation mechanisms in a distributed architecture.

Timing of mechanism	*Location of mechanism*	*Kind of minimality applicable*	*Privacy property targeted*
before deployment	clients	distributed	collection minimisation
	3rd party	monolithic	
	server		use minimisation
online	clients	distributed	collection minimisation
	3rd party	monolithic	
	server		use minimisation
offline	clients	distributed	breach detection
	3rd party	monolithic	
	server		

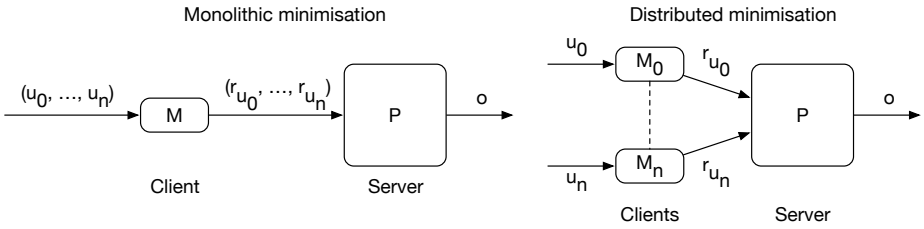

Fig. 2: Monolithic and distributed minimisation architectures.

by the server, then monolithic minimality is achieved as there is full knowledge about input values.

Privacy Properties The kind of privacy property that can be achieved depends on when in the data lifecycle the methods are deployed. *Collection minimisation* prevents data from being collected by the server (this is the case when the client itself, or a trusted third-party, performs the minimisation). On the other hand, *use minimisation* is only able to limit the use of the data (this is the case when the server itself is also in charge of performing the minimisation[6]).

We note that the only way to implement collection minimisation while ensuring monolithic minimality in a distributed setting is by using a third party to perform the minimisation. Indeed, it is not possible to reach a property stronger than distributed minimality when the minimisation is applied client-side.

Figure 2 gives an illustration of monolithic vs. distributed minimisation in the case of a static minimiser generated before deployment, and run by the client (thus guaranteeing collection minimisation).

In the following, we detail the attacker model used to formally characterise the notion of minimality. In the following, unless otherwise stated, will focus on the static mechanism built before deployment for verification purposes; this

[6] One could argue that it is already too late in this case, but the server could belong to a multi-server organisation, thus minimising potential subsequent disclosures (not treated in this paper).

is the most complex case, and the other mechanisms can be seen as weakened versions of this one.

3 Attacker Model

In order to define data minimisation we have to consider an explicit attacker model. For example, if the computation performed on the data is really just the intended program P, then why worry about minimisation? In practice a malicious[7] server may have a *secondary use* for the data (defined as "using personal data in a different way than the primary reason why they were collected" [7]). We define here an attacker model having the following components: (i) an explicit model of a hidden secondary use of the data, and (ii) an explicit model of what the attacker learns about the data from observing the result of processing the data. We will develop a definition of minimisation which guarantees that an attacker learns no more about the private data than he can deduce from the legitimate output of the program.

To model the hidden secondary use, we suppose the attacker uses a second program denoted by a function $h \in \mathcal{I} \to \mathcal{O}'$. Thus the attacker is defined by a pair of functions p and h (the legitimate use and the hidden use of the data, respectively), and the attacker's computation is given by the function $\langle p, h\rangle \in \mathcal{I} \to \mathcal{O} \times \mathcal{O}'$ defined by

$$\langle p, h\rangle(i) \stackrel{\text{def}}{=} (p(i), h(i))$$

The goal of the attacker is to learn something about the inputs by observing the output of $\langle p, h\rangle$. In the following section we will show that if the input is first minimised using a best possible minimiser, then what the attacker learns from $\langle p, h\rangle$ is no more than what is learned by p alone. To do this we model the attacker knowledge explicitly, as has become more common in recent formulations of the semantics of information flow (e.g. [3, 4, 6]).

Definition 1 (Knowledge). *Let $u \in \mathcal{I}$ be an input, and f a function in $\mathcal{I} \to \mathcal{O}$. We say the* knowledge *about u obtained through f by observing $f(u)$, written $K(f, u)$, is defined to be the set $\{v \mid f(u) = f(v), v \in \mathcal{I}\}$.*

Thus $K(f, u)$ is the largest set of possible input values that could have led to the output value $f(u)$. This corresponds to an attacker who knows the function f and the output $f(u)$ and makes perfect logical deduction about the possibles values of the input. Note that the smaller the knowledge set, the less uncertainty the observer has over the possible value of the input.

[7] In security literature, a malicious attacker is an attacker which does not follow the protocol while an attacker only attempting at extracting more information from what he got is called honest-but-curious. In the privacy context, the (security) honest-but-curious attacker cannot be considered as honest because personal data should not be used for another purpose than the one specified and consented for by the data subject. As a consequence, this is also considered as a malicious behaviour.

The knowledge relation induces an order on programs: the bigger the knowledge set for a given input, the lesser a program will disclose information about the input. This is reasonable given that if more input values are mapped to the same output, more difficult it will be to guess what is the exact input when observing the output of a computation. We have then the following order between programs:

Definition 2 (Disclosure ordering). *We say* f discloses less or as much information as g, *written* $f \sqsubseteq g$, *iff for any* $u \in \mathcal{I}$ *we have* $K(g,u) \subseteq K(f,u)$.

If f and g are knowledge equivalent, we write $f \equiv g$. For example, functions denoted by $(\cdot \bmod 2)$ and $(\cdot \bmod 4)$ (with postfix notation) are related by $(\cdot \bmod 2) \sqsubseteq (\cdot \bmod 4)$ (knowing the last bit discloses less information than knowing the last two bits). However, $(\cdot > 0)$ is incomparable to either of them.

4 Data Minimisers

In this section we develop a definition of a *minimiser* for a program. We will consider two types of minimiser depending on how the data is collected: *monolithic* and *distributed*. Monolithic collection describes the simple case when the personal data is collected from a single location (e.g. a single data subject). On the other hand, distributed collection occurs when personal data is collected from different sources. In the latter case we assume these sources do not interact to achieve a multi-party data minimisation. This assumption is important as it has a strong impact to the kind of minimisers which may be feasible in practice.

The execution of a program consists in the computation of an output, depending on the inputs provided. The domain of possible inputs can be arbitrarily large and complex. Running a data minimiser before a program is intended to capture the idea that the program will be given less data. In the best case, only the data actually needed to produce the result will be provided. Thus it implies that the system does not receive more sensitive data than it needs.

We will address the whole data minimisation process by the following steps: (i) data minimiser generation (from the program, performed client-, server-, or third-party-side), (ii) data minimisation *per se* (ie. data minimiser execution, performed client-side), (iii) data collection (ie. representative transmission to the server, performed client-side), and (iv) program execution (performed server-side). Note that a minimiser only needs to be generated only once for a given program: the same can be used for all clients. This section and the following one are mainly concerned with the first of these steps (the other steps being simple executions and communications). The improvement in the minimality of a given program is achieved by constructing a *minimiser* M.

In the following we assume the data processor to be the malicious attacker depicted in Section 3 and the data sources to be as trustable as they would be without data minimisation. The data minimiser could be sent as part of the protocol for data collection. This may be done by a third-party or by the data

```
1 input(x1, x2, x3);
2 y := 0;
3 if (x2 == x2;)
4     then (y := (x3 + x1 - x3);)
5 output(y);
```

Fig. 3: Example of a semantically unnecessary syntactic necessity.

```
1 input(salary)
2 repr_salary := 0
3 if (10000 <= salary)
4     then (repr_salary := 10000)
5 output(repr_salary)
```

Fig. 4: Minimiser for P_{bl} (see Figure 1).

processor itself. Thus we assume that the data source knows how much data is minimised on its side, but may not be able to check the optimality of the whole minimiser. A third-party provider for the data minimiser could help to get this assurance. Note, however, that the fact that the specification is being determined by the data processor itself does not weaken (nor strengthen) the security and trust assumptions of the interactions between the parties.

4.1 Minimisation Semantics

Let us consider the program in Fig. 3 to discuss the need for a semantic definition of minimisation. In this program, `x2` is syntactically needed to evaluate the condition (l.3). However, this condition always evaluates to `true`. In the same way, `x3` is not semantically needed to compute the value of `y` since it is both added and substracted (l.4). As a consequence, it would be possible to get the same result by rewriting the program with only the input `x1` without modifying its behavior. If `x2` and `x3` are personal data, then the semantic approach is better likely to help us to limit the collection than the syntactic approach. So, the program could be refactored by taking only the input `x1` while retaining the same output behaviour. Though this would work, this approach requires a change in both the data collected and in the interface of the program.

Instead, we propose to keep the program unchanged and we rely on the idea that the information behind `x2` and `x3` (in this specific example) can be protected by providing *fixed arbitrary* values for them, instead of refactoring the program. This means the program does not need to be modified for the data processor to propose better guarantees. This approach allows a better modularity and a better integration in the software development cycle. To see this, let us consider the program shown in Fig. 1. In this case, any employee earning less than $10000 can disclose any figure between 0 and 9999 and any employee earning at least $10000 can disclose any figure between 10000 and 100000 without affecting the output of the program. Thus a corresponding *data minimiser* could be as shown in Fig. 4, where the representative values are taken to be `0` and `10000`.

The value output by this minimiser, which is the value of `repr_salary`, can then be provided as input to the corresponding program (standing for the variable `salary`). The behaviour of the program will remain the same while the actual salary of the employee is not needlessly passed to the program. This holds even when an employee earns exactly the amount of `repr_salary` since the data processor receiving this value cannot distinguish when it is the case or not.

Following the approach about data minimality, we study in turn both monolithic and distributed cases as shown in Fig. 2 where M and $\bigotimes_{i=0}^{n} M_i$ are a monolithic and a distributed minimiser respectively (each M_i being a local minimiser), P is the program, u_i are input values, r_{u_i} are input representatives, and o are outputs. We then discuss about security properties achieved through composition of minimisers and programs.

4.2 Monolithic Case

Let us assume a program P constituting a legitimate processing of sensitive (personal) data for which the data subject had consented to the purpose. We abstract it through a function p which ranges over $\mathcal{I} \to \mathcal{O}$. Since we aim at building *pre-processors* for the program, we consider a minimiser m with type $\mathcal{I} \to \mathcal{I}$.

As stated previously, monolithic collection refers to the case when there is only one data source from which data is collected before being processed. This data source can thus minimise the data disclosed while having full knowledge about the input that are provided to the processor. Thus the output of a data minimiser M can be directly input to the program P.

Definition 3 (Monolithic minimiser). *We say m is a* monolithic minimiser *for p iff (i) $p \circ m = p$ and (ii) $m \circ m = m$.*

Condition (i) on the above definition is basically correctness: it states that using the pre-processor does not change the result of the computation. Condition (ii) ensures that m chooses representative inputs in a canonical way. Put it another way, m has the job of removing information from the input. This idempotency condition states that m removes all the information in one go.

4.3 Distributed Case

In general a computation over private data may require data collected from several independent sources. We will thus extend our definition to the distributed setting. This is the general case where a program dp is a function of a product of input domains such that $dp \in \prod_{i=0}^{n} \mathcal{I}_i \to \mathcal{O}$. The idea for the distributed minimiser will be to have a *local minimiser* $m_i \in \mathcal{I}_i \to \mathcal{I}_i$ for each source $\mathcal{I}_i$, combined into an input processor as $dm = \bigotimes_{i=0}^{n} m_i$, where for $f \in A \to A'$ and $g \in B \to B'$ we have $f \otimes g \in A \times B \to A' \times B'$. This is based on the assumption that each argument of dp is provided by a different data source and hence requires its own data processing.

Definition 4 (Distributed minimiser). *We say dm is a* distributed minimiser *for dp iff (i) dm is a monolithic minimiser for dp and (ii) dm has the form* $\bigotimes_{i=0}^{n} dm_i$ *(with* dm_i *being* $\mathcal{I}_i \to \mathcal{I}_i$ *functions (called local minimisers, for* $i \in \{0, \ldots, n\}$*)).*

The first condition ensures that *dm* is actually a minimiser while the second condition (ii) ensures that each input is treated independently. This second condition is the key differentiator between monolithic and distributed minimisers. Each function dm_i is only provided with an input from $\mathcal{I}_i$. Intuitively, this means that if two input values at a given position belong to the same equivalence class, then it has to be the case that this holds for all values possibly provided for the other positions. For example, if the values 0 and 1 from $\mathcal{I}_0$ belong to the same equivalence class, then each pair of input tuples $\langle 0, x, y\rangle$ and $\langle 1, x, y\rangle$ for any $(x, y) \in \prod_{i=1}^{2} \mathcal{I}_i$ have to belong to the same equivalence class (two by two, meaning these can be different classes for each pair: $\langle 0, 1, 2\rangle$ and $\langle 1, 1, 3\rangle$ may not belong to the same equivalence class, though $\langle 0, 1, 2\rangle$ and $\langle 1, 1, 2\rangle$ have to). This is formally stated in Proposition 1 below, which relies on the definition of the *kernel* of a function.

Definition 5 (Kernel of a function). *If* $f \in \mathcal{I} \to \mathcal{O}$ *then the* kernel of f *is defined as* $\ker(f) = \{(u, v) \mid f(u) = f(v)\}$.

So, the kernel $\ker(f)$ induces a partition of the input set where every element in the same equivalence class is mapped to the same output.

Proposition 1. *If m is a monolithic minimiser for p, then m is a distributed minimiser for p iff for all* $(\boldsymbol{u}, \boldsymbol{v}) \in \ker(m)$*, for all input positions i, and all input vectors* $\boldsymbol{w}$*,*

$$(\boldsymbol{w}[i \mapsto u_i], \boldsymbol{w}[i \mapsto u_i]) \in \ker(m)$$

where the notation $\boldsymbol{w}[i \mapsto u_i]$ *denotes a vector like* $\boldsymbol{w}$ *except at position i where it has value* u_i*.*

This proposition gives a data-based characterisation of *data minimisation*. This will be useful when building minimisers in Section 6. Before building minimisers, we explain how to compare them in the next section.

5 Best Minimisers

Now that we defined minimisers as pre-processors modifying the input, we see that there may exist many different minimisers for a given program. Thus we are interested in being able to compare these minimisers. Indeed, since the identity function is a minimiser for any program p, then it is clear that the simple existence of a minimiser does not guarantee any kind of minimality. One way to compare minimisers is to compare the size of their ranges – a smaller range indicates a greater degree of minimality (cf. Proposition 2 below). But a more precise way to compare them is by understanding them in terms of the *lattice of*

equivalence relations [26], which has been used to model dependency properties of programs in several works, e.g. [17, 19, 28].

The set of equivalence relations on $\mathcal{I}$ forms a complete lattice, with the ordering relation given by set-inclusion of their defining sets of pairs. The identity relation (which we denote $\mathrm{Id}_{\mathcal{I}}$) is the bottom element, and the total relation (which we denote $\mathrm{All}_{\mathcal{I}}$) is the top.

The following proposition provides some properties about the order relation between programs (cf. Definition 2), including its relation with the kernel.

Proposition 2 (Disclosure ordering properties).
1. $f \sqsubseteq g$ *iff* $\ker(g) \subseteq \ker(f)$
2. $f \circ g \sqsubseteq g$
3. $\langle f, f \rangle \sqsubseteq f$
4. $f \sqsubseteq \langle f, g \rangle$
5. $f \sqsubseteq g$ *implies* $|\mathrm{range}(f)| \leq |\mathrm{range}(g)|$

where the notation $\langle f, g \rangle$ a function denoting a mapping of an argument to each one of f and g as defined in Section 3.

5.1 Monolithic Case

In order to reason about minimisers we need to be able to compare them. The ordering relation between programs defined above may be specialised to minimisers: m is a *better minimiser* for p than n is iff (i) m and n are minimisers for p and (ii) $m \sqsubseteq n$. Of particular interest is the minimiser that fulfill its purpose in the *best* way:

Definition 6 (Best monolithic minimiser). *We say that m is a* best monolithic minimiser *for p iff (i) m is a monolithic minimiser for p and (ii) $m \sqsubseteq n$ for any n a minimiser for p.*

A minimiser carrying less information than the best minimiser could not convey all the information necessary for the program p to keep the same behaviour under all inputs. In this simple (monolithic) form, minimisation is easily understood as injectivity:

Proposition 3 (Monolithic best minimiser injectivity). *A monolithic minimiser m for a program p is a best one iff $p|_{\mathrm{range}(m)}$ is injective (with $p|_{\mathrm{range}(m)}$ the restriction of the program p over the range of m).*

Now we can show that using a minimiser m at the client guarantees that the attacker $\langle p, h \rangle$ learns no more about the input than which can be learned by observing the output of the legitimate use p (the proof follows from the ordering between minimisers and Proposition 2):

Theorem 1. *If m is a best monolithic minimiser for p then for any hidden use h we have $\langle p, h \rangle \circ m \equiv p$.*

Proof.

$$
\begin{array}{lll}
p & \equiv p \circ m & (m \textit{ is a minimiser for } p) \\
& \sqsubseteq \langle p \circ m, h \circ m \rangle & (\textit{Prop. 2(4)}) \\
& \equiv \langle p, h \rangle \circ m & (=)
\end{array}
$$

$$
\begin{array}{lll}
\langle p, h \rangle \circ m \sqsubseteq m & & (m \textit{ is a minimiser for } p, \\
& & \textit{Prop. 2(2)}) \\
& \sqsubseteq \langle p, m \rangle & (\textit{Prop. 2(4)}) \\
& \sqsubseteq \langle p, p \rangle & (\textit{Def. 2}) \\
& \sqsubseteq p & (\textit{Prop. 2(3)})
\end{array}
$$

Finally, we note that:

Theorem 2. *For every program p there exists a best monolithic minimiser.*

5.2 Distributed Case

As for the monolithic case, we have an ordering between distributed minimisers. The notion of a best minimiser also exists in the distributed case and restricts the monolithic case definition as follows:

Definition 7 (Best distributed minimiser). *We say dm is a* best distributed minimiser *for dp iff (i) dm is a distributed minimiser for dp and (ii) $dm \sqsubseteq dn$ for any dn a distributed minimiser for dp.*

The following result shows that there always exists a best distributed minimiser (see proof in Appendix A).

Theorem 3. *For every distributed program dp there exists a best distributed minimiser.*

Note, however, that there is a price to be paid for having distributed minimisers: the best distributed minimiser may reveal more information than the best monolithic minimiser for a given function. An example is the Boolean *OR* function where the identity function is a best distributed minimiser, but in the monolithic case we can minimise further by mapping $(0, 0)$ to itself, and mapping all other inputs to $(1, 1)$.

Similarly to what we did in the previous section with Proposition 1, we give a data-based characterisation of best minimisers as follows.

Proposition 4 (Data-based best distributed minimisation). *If dm is a best distributed minimiser for dp, then for all input positions i, for all v_1 and $v_2 \in \mathcal{I}_i$ such that $v_1 \neq v_2$, there is some $\boldsymbol{u} \in \text{range}(dm)$ such that*

$$dp\left(\boldsymbol{u}[i \mapsto v_1]\right) \neq dp\left(\boldsymbol{u}[i \mapsto v_2]\right).$$

This property will prove useful in the next section to build minimisers.

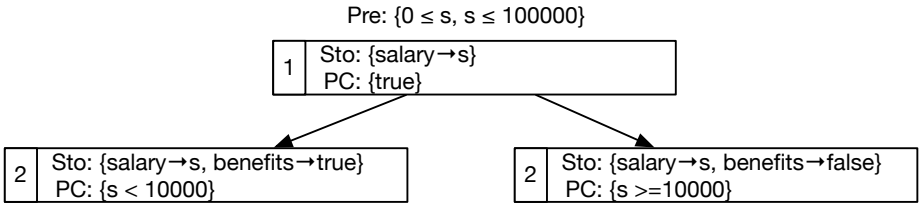

Fig. 5: Symbolic execution tree of P_{re} (see Figure 1).

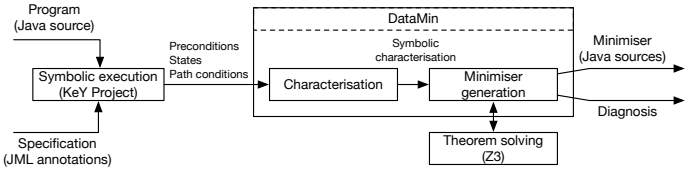

Fig. 6: Toolchain of the (semi-)procedure to generate minimisers.

6 Building Minimisers

We describe here a procedure to build data minimisers. This (semi-) procedure is not complete for the best minimisers since obtaining a best minimiser is not computable in general. Besides, and more pragmatically, our procedure is built on top of a theorem prover and a symbolic execution engine, and it is thus also limited by the power of such tools in addition to the theoretical algorithmic limits.

Only the main program (and optionally a specification overriding the implicit specification embodied in the program itself) is needed to generate a minimiser. Data minimisers thus generated are intended to be deployed either monolithically or in a distributed manner with a part residing at each data collection point; each data source locally executes the local minimising program before forwarding the representative inputs to the data processor (i.e. the program).

Let us consider again the example described in Fig. 1 to give an intuition of the way minimisers can be generated. In order to generate a data best minimiser for the program we need to reason about the dependencies between the possible outputs and the inputs considering all possible execution paths. A representation of the possible paths for this program is shown in Fig. 5 where the digits 1 and 2 correspond to the lines of the program, Sto is the corresponding store, and PC is the condition to be satisfied at this point (l.3 of the program does not have any effect in this representation).

The possible execution paths for this program lead to two main outputs, where either `benefits == true` or `benefits == false`. Thus the best minimiser generated should distinguish the two cases and return a representative value leading to the same output as if the program were to be executed with the "real" value assigned to `salary`. This is the case for the data minimiser described in Fig. 4. Moreover, we see that any change of value for `repr_salary` leads to a change in the `salary` computed by the main program. This means this minimiser is indeed a best minimiser.

The procedure to generate data best minimisers follows the toolchain depicted in Fig. 6. Before describing the procedure in detail, we briefly recall concepts related to *symbolic execution*, since we run the program symbolically in order to generate the best minimiser.

6.1 Symbolic Execution

This step corresponds to the two leftmost boxes of the toolchain shown in Fig. 6. Symbolic execution has been pioneered in [18] as a way to reason about symbolic states reached by a program execution.

A *symbolic execution* of a program is an execution of the program against symbolic inputs. Expressions are symbolically evaluated and assigned to variables. The state of this execution includes the values of the program (modelled as a mapping called *store*), and *path conditions* (boolean expressions ranging over the symbolic inputs, specifying the conditions to be satisfied for a node to be reached). Assignments modify the symbolic value of variables, whereas conditions and loops create branches distinguished by their path conditions.

This symbolic execution generates a *symbolic execution tree* where each node is associated to a statement and each arc is associated to a transition of the program. A state is associated to each node and the root node corresponds to the input of the program. Each leaf in the tree then corresponds to a completed execution path, and no two leaves have the same path condition.

A key condition for the soundness of symbolic execution is the commutativity of its semantics with respect to the concrete semantics of the language. This means that executing the program symbolically on given symbolic values and then concretising the output should give the same result as executing the original program over the concretisation of the symbolic inputs.

The formal definition of this tree makes it possible to reason about the program under analysis. However, some executions may lead to infinite trees in the presence of loops (when they cannot be simply unfolded). In this case, in order to get a procedure that terminates and provide conditions for getting a symbolic execution of the program, one should provide *loop invariants* for each concerned loop in the program. To find the best invariant is a difficult problem and still an active field of research [12]. Failing to provide the best invariant may result in weaker path conditions failing to capture all of what is theoretically knowable about the states of the symbolic execution.

In our approach the program P is symbolically executed along with the assertions coming from the specification S. We assume that we have a global *precondition* for the program P, given in the specification S of the program, denoted $\mathrm{Pre}\langle P, S\rangle$. This specification can be expressed explicitly (with annotations in the source code of P for example), or implicitly (by e.g. analyzing the range of the program during the symbolic execution). In our example, the conditions $0 \leq s$ and $s \leq 100000$ are part of the preconditions as shown in Fig. 5. This execution produces a symbolic execution tree equipped with the corresponding path conditions PC and Sto attached to its nodes.

In what follows we define a symbolic characterisation of the program P under specification S capturing the conditions at the termination of the execution.

Definition 8 (Symbolic characterisation of a program). *We say Γ is a* symbolic characterisation *of a program P under specification S, written $\Gamma_{\langle P,S\rangle}$, iff Γ collects the preconditions, stores, and path conditions to be satisfied for each possible output of P: $\Gamma_{\langle P,S\rangle} \triangleq \mathrm{Pre}\langle P,S\rangle \wedge \Big(\bigvee_{l\in \mathrm{Leaves}\langle P,S\rangle} (\mathrm{PC}(l) \wedge \mathrm{Sto}(l))\Big)$, where $\mathrm{Leaves}\langle P,S\rangle$ returns the leaves of the symbolic execution tree of P under specification S, and $\mathrm{PC}(\cdot)$ and $\mathrm{Sto}(\cdot)$ return the path condition and the state associated to a leaf, respectively.*

For the example in Fig. 1, the (simplified) symbolic characterisation is: $\Gamma_{\langle P_{bl},S_{bl}\rangle} = (0 \leq s \wedge s \leq 100000) \ \wedge \ ((s < 10000 \wedge \mathit{salary} = s \wedge \mathit{benefits} = \mathit{true}) \vee (s \geq 10000 \wedge \mathit{salary} = s \wedge \mathit{benefits} = \mathit{false}))$

The symbolic characterisation induces an equivalence class giving a partition of the state space according to the different possible outputs. This is then used by a solver to get a representative for each equivalence class, which is the next step in our procedure.

6.2 Static Generation Before Deployment

In this section we show how best minimisers can be generated from the symbolic execution tree. In the following, the monolithic case is seen as a specific case of the distributed case: the former is reduced to a distributed case with only one input. The construction spans over the two boxes of the toolchain in Fig. 6 as detailed in the following paragraphs.

The input domain of the semantic function *dp* having $n+1$ inputs is $\prod_{i=0}^{n} \mathcal{I}_i$ and the corresponding input variables of *dp* are denoted as x_i. Each domain $\mathcal{I}_i$ is determined implicitly by *DP* or explicitly by its specification S. A local function dm_i is generated for each input $x_i \in \mathrm{Input}(DP)$. We are still interested in the different possible outputs of *dp* but we cannot directly use its kernel as in the monolithic case since this would require all the minimisers to have access to the data at other data source points. Instead of this, each data source point assumes that the data to be disclosed by other points can take any value in their domain. Thus we need to compute, for each input variable x_i, the equivalence classes over its domain such that the value o of the output of the program remains the same for any possible assignment to the other inputs variables x_j for any $j \in (\{0,\ldots,n\} \setminus \{i\})$.

The algorithm of Fig. 7 shows how the distribution of the inputs is taken into account. Here, *min* (l.1) stands for the minimiser being built (as a map containing the local minimiser *min*[i] (l.5) for each input i (l.3)), *Inputs* (l.3) denotes the inputs of the program, and Γ (l.4) denotes the formula of the symbolic characterisation of the program. The notation $\phi[y/x]$ denotes the formula ϕ in which y replaces x and == the logical equality (l.8) while $S \setminus s$ denotes $S \setminus \{s\}$ (the set S from which the element s has been substracted, l.9).

1: $min = \{\}$ ▷ initialise minimiser
2: decl logical_variable i' ▷ declare a new logical variable
3: **for all** $i \in Inputs$ **do** ▷ iterate over the inputs
4: $\gamma = \Gamma$ ▷ copy the program symbolic characterisation to γ
5: $min[i] = \{\}$ ▷ initialise minimiser for input i
6: **while** γ.check() **do** ▷ **call solver** to loop as long as γ is satisfiable
7: $model = \gamma$.model() ▷ **call solver** to get a valid assignment for γ
8: $formula = i' == model[i] \wedge \Gamma \wedge \Gamma[i'/i]$ ▷ build distributed minimisation formula
9: $formula = \exists i'.\forall (Inputs \setminus i) \in Dom(Inputs \setminus i).\exists o.formula$ ▷ introduce quantifiers except for i
10: $wp = formula$.quantif_elim() ▷ **call solver** to eliminate quantifiers
11: $min[i] += (wp, model[i])$ ▷ add weakest precondition with representative to minimiser
12: $\gamma = \gamma \wedge \neg wp$ ▷ conjunct negation of the weakest precondition to γ to limit the loop
13: **end while**
14: **end for**

Fig. 7: Distributed data minimiser generation (excerpt).

Primitives depending on a theorem prover are called at three locations in the algorithm. The first one, check() (l.6) checks whether or not a logical formula is satisfiable. Then, model() (l.7) applied to a formula returns a satisfying valuation. These two primitives are linked as it is only possible to find such a valuation for satisfiable formula. Finally, quantif_elim() (l.10) is a procedure to eliminate (both universal and existential) quantifiers in a formula, thus simplifying a formula by removing the corresponding symbols.

After having initialised the data structure *min* holding the minimisers as a map from inputs i to tuples (*weakest precondition*, *representative*) (l.1), a new logical variable i' is declared (l.2). At this point, all the inputs (from *Inputs* (l.3)) and the output o of the program already exist as logical variables (implicitly introduced with Γ which is an input of this algorithm). The rationale to introduce a new logical variable i' is because it is used to control variations on one input during the procedure (l.8-9). Then, the algorithm iterates over all the inputs (l.3). The symbolic characterisation Γ is assigned to the variable γ (l.4) (which will be reinitialised at each iteration of the loop of l.3). The original Γ will be used to build formulas to be solved while the fresher γ will be used to control the loop of l.6. The minimiser *min* is then initialised for the current input i (l.5). The algorithm should then loop over all equivalence classes for the current input i. This is ensured by (i) representing each equivalence class by its weakest precondition *wp* (l.7-11), (ii) conjuncting the negation of the weakest precondition *wp* found to the symbolic characterisation γ (l.11), and (iii) looping as long as there is another equivalence class (ie, weakest precondition) to be covered by checking the satisfiability of γ conjuncted with the negation of all previous conditions (l.6).

We now explain in more detail how the weakest preconditions are found (l.7-11). A satisfying valuation of the characterisation γ is requested from the solver and assigned to the variable *model* (l.7). The valuation of variable x can be called by $model[x]$. A formula (assigned to the variable *formula*) is then built in two steps. First, we conjunct Γ, $\Gamma[i'/i]$, and the formula $i' == model[i]$ (l.8). This fixes $\Gamma[i'/i]$ to the case where i' is equal to the satisfying valuation found previously (l.7). Then we introduces quantifiers for all the free logical variables except for i to get the data-based characterisation of data minimisation as given by Proposition 1 (l.9). The output o is existentially quantified as it can freely vary without any other constraint than those imposed by the Γs. All the inputs except for i and i' are universally quantified because the characterisation states that the preservation of the functional correctness (implied by the value of the output o here) has to be guaranteed for all possible variations of inputs (except for inputs i and i'). Finally, the input i' is existentially quantified only for the purpose of being eliminated later (free variables are existentially quantified for SAT solvers). Once the formula has been built, the quantifiers are eliminated by calling the suitable procedure from a solver (l.10). This gives the weakest precondition corresponding to an equivalence class of the inputs (the one corresponding to the value of $model[i]$). This equivalence class *wp* is then added to the minimiser $min[i]$ along with its representative $model[i]$ (l.11) before being excluded from the control γ (l.12) before a new iteration, until exhaustion of equivalence classes (l.6).

Theorem 4 (Soundness). *The algorithm of Fig. 7 builds a best distributed minimiser dm for program dp.*

The soundness of this algorithm relies on Proposition 4, on the fact that the representative assigned to each equivalence class is fixed, and on the proof of existence in Appendix A.

Theorem 5 (Termination). *The algorithm of Fig. 7 terminates when the number of inputs, the number of equivalence classes on these inputs, and the calls to the external solver terminate.*

This is proven by showing that all the loops in the algorithm iterate over sets built from inputs and equivalence classes. The set of inputs is not modified in the loops while the condition of (l.6) strictly decreases in terms of the number of times it is satisfiable. However, it terminates only if there is a finite number of equivalence classes for the current input. Since we depend on external procedures, the termination condition also rely on these ones satisfying this condition. In practice, this is only a semi-procedure.

In both the monolithic and the distributed cases, generating a minimiser may give interesting insights about the process. It can be used to get metrics, such as the number of equivalence classes finally obtained, to compare with the initial size of the input domain. It could also be used for quantitative (information-theoretic) analysis.

6.3 Other Mechanisms

The mechanism presented in the previous section is the most expensive to deploy. Indeed, the static approach requires computation of all the (symbolically) possible cases. Thus, each of the calls to the external solver in the algorithm depicted Fig. 7 occurs $\sum_{i \in Inputs} \#equivClasses_i$ times with $\#equivClasses_j$ being the number of equivalence classes for the input i (there is in fact one more call for the satisfiability check (l.6)). This cost is hard to predict in advance. Indeed, programs which can be highly minimised and end up with few equivalence classes can offer good performance, but quasi-injective programs may have a large number of equivalence classes.

This complexity justifies the need of other mechanisms such as the ones presented in Table 1. All these other mechanisms allow the reduction of the complexity by limiting the domain of verification. To effectively use data minimisation, it is thus necessary in general to choose easier mechanisms. For instance, the static and dynamic mechanisms run online after deployment leads to only $\sum_{i \in Inputs} 1 = |Inputs|$ calls to the external procedures since the input is known at run-time. However, this still needs to call semi-decidable external procedures, which may create a large overhead. If weaker guarantees are acceptable, then it is possible to use test strategies which verify only partial domains.

7 Implementation and Examples

In this section we present *DataMin*, an implementation of the procedure to generate a minimiser. Two examples are then given to illustrate the notions of data minimality and data minimisation. The first example contains a loop and lies in the monolithic case while the second example focuses on a distributed setting to illustrate the distributed minimisation.

7.1 *DataMin* Implementation

The (semi-)procedure described in the previous section has been implemented in the Python language as a proof of concept named *DataMin*.We rely on the Symbolic Execution Engine of the *KeY* Project [16]. This symbolic executor is run against a program P written in Java source code for which a minimiser should be generated.

The program can be instrumented with Java Modeling Language (JML) annotations [20] to help *KeY* with the exploration of the states space. They can be used to give a specification S of the domain of inputs under normal behavior for the program according to the design-by-contract terminology [23] (written `//@ requires P` with `P` a property assumed to be verified). The annotations can also be helpful when dealing with loops that cannot be unfolded. In this case, a loop invariant can be provided (written `//@ maintaining P`, with `P` the loop invariant) along with a description of a decreasing term to ensure termination (written `//@ decreasing t`, with `t` the decreasing term).

The symbolic execution results in the definition of stores and path conditions for each leaf of the symbolic execution tree along with the preconditions when applicable. *DataMin* generates the symbolic characterisation $\Gamma_{\langle P,S\rangle}$ and builds the partitioning $k_{\langle P,S\rangle}$ and the sectioning $r_{\langle P,S\rangle}$ functions. The theorem prover *Z3* [24] is called through its API to solve constrains as needed during these steps. This proof-of-concept supports only a limited range of data structures but could be extended thanks to the many theories on which *Z3* relies.

Finally, *DataMin* generates the minimiser as a set of Java source code files to be run by the data source points before disclosing the data. These files are directly compilable and ready to be exported as Java libraries to ease their use. This whole process currently runs in reasonable time for these examples (less than a second, given that Python is run in its interpreted mode). For the first example, the solver used (*Z3*) is not able to totally perform quantification elimination. We thus checked the results hold by manually verifying them using the *Redlog* system [10]. This is a limitation of the external tools and not of the procedure proposed in this paper.

7.2 Loyalty Status

The first example proposes a situation where an airport facility must deliver different services to different customers depending on their status. This status level is determined depending on the number of flights made by the customer in the previous year with its favorite company, *PrivaFly*. This number of flights is disclosed by the company to the airport.

However, *PrivaFly* wants to adopt the best practices in personal data protection and requires that only the needed data should be disclosed. The airport services have their own policy to compute the status level, as shown in Fig. 8. The method `compStatusLevel` takes a number of `flights` and returns a `status` level. If the number of flights is lower than 9, the status is 0. From 10 to 19 flights, the status level is equal to the number of flights from which 10 is substracted. For a number of flights ranging from 20 to 29, the computation involves a loop and will be detailed later. Finally, over 30 flights, the status level is capped to 500. The first JML annotation (l.2) specifies a range for the input domain. Then, the following annotations from (l.9-13) will be described later.

Intuitively, there is no need to give a precise value for a number of flights between 0 and 9 and over 30. On the other hand, the exact number should be disclosed between 10 and 19. The program M_{ls} from Fig. 9 is the result of running *DataMin* on P_{ls} (see Fig. 8). In this program, the complex (weakest pre)conditions for each equivalence class returned by *DataMin* have been replaced by placeholders `wp`i (with i a number). Each of these conditions depends only on the input of the minimiser (here, on the variable `flights`) in order to be executable.

We can see that the intuition above is partially confirmed by the minimiser generated. The representatives chosen by the minimiser for `repr_flights` are the minimum of the corresponding equivalence classes. As a consequence it has been fixed to 0 when the actual number of flights ranges from 0 to 10 (l.4 of

```
1  public class LoyaltyApp {
2      //@ requires (0 <= flights) && (flights <= 100);
3      public int compStatusLevel(int flights) {
4          int status;
5          if (flights < 10) { status = 0; }
6          else if (flights < 20) { status = (flights - 10); }
7          else if (flights < 30) {
8              status = 0; int i = 0;
9              //@ maintaining (0 <= i)
10             //@            && (i <= flights - 19)
11             //@            && (status == i * flights);
12             //@ decreasing (flights - 19 - i);
13             //@ assignable \strictly_nothing;
14             while (i <= flights - 20) {
15                 status = (status + flights);
16                 i = i + 1; }
17             if (status > 150) { status = 150; } }
18         else { status = 500; }
19         return status; } }
```

Fig. 8: Program P_{ls} to compute a loyalty status.

```
1  public class LoyaltyAppMin_flights {
2      public int minimise_compStatusLevel(int flights) {
3          int repr_flights;
4          if (wp1) { repr_flights = 0; }
5          else if (wp2) { repr_flights = 11; }
           ...
18         else if (wp15) { repr_flights = 24; }
19         else if (wp16) { repr_flights = 25; }
20         else { repr_flights = 30; }
21         return repr_flights; } }
```

Fig. 9: Minimiser M_{ls} for P_{ls} (see Fig. 8).

the minimiser). We observe 10 is included (case not captured by the following conditionals of the minimiser) despite the fact it is not captured by the first conditional of the program (l.5) but by the second one (l.6). However, since the status level is also equal to 0 in this case, it belongs to the same equivalence class as the number of flights captured by the first conditional. For a number of flights between 11 and 19, the output is different each time (number of flights minus 10). Thus the representative is the number of flights itself (l.5-13 of the minimiser). The case of a number of flights greater than 30 is trivially captured by a representative fixed to 30 as well (l.20).

Between 20 and 29 flights, the status is computed thanks to a loop (l.14-16 of the program). The code makes the status level computed after this point to be equal to the number of flights multiplied by the same number minus 19. The symbolic executor of *KeY* is not able to solve this loop itself and an invariant (`maintaining`) has been provided (l.9-11). These relations described by

Table 3: Equivalence classes for P_{ls} (see Fig. 8).

invariant	*variables* \ *values*	$[0,9]$	10	$[11,19]$	$[20,24]$	$[25,29]$	$[30,100]$
	`flights`	1^{st} `if`		2^{nd} `if`	3^{rd} `if`		final `else`
best	`repr_flights`	1 class		9 classes	5 classes	1 class	1 class
weak	`repr_flights`	1 class		9 classes	5 classes	5 classes	1 class

Note: Correspondance between `flights` and their representatives, mismatch between syntactic structure (`if` conditionals) and semantic classes, and effects of weak invariants on the minimisation.

the loop invariant are true both at the beginning and at the end of each iteration. The following annotation (`decreasing`, l.12) helps the executor with the termination of the loop by providing a decreasing term. Finally, the last annotation (`assignable`, l.13) also helps the symbolic executor to handle the loop (by declaring here that no side effect occurs). These annotations give to the symbolic executor the best invariant. As a consequence, the data minimisation process is able to propose a minimiser. The representatives for a number of flights from 20 to 24 correspond to the number of flights itself (l.14-18 of the minimiser). This is not the case starting from 25 until 29 (l.19 of the minimiser) because the status is capped at 150 when the number of flights is less than 30 and these five possibilities constitute their own equivalence class.

If the invariant provided for the loop had not been the best invariant, the equivalence class composed of a number of flights between 25 and 29 may not have been detected. The representatives for these cases would have been the actual number, leading to more disclosure than needed.

The observations made during this example are summed up in Table 3. The table shows the mismatch between the syntactic structure of a program (modelled in this table by the correspondence between the number of flights and the `if` conditionals which are triggered) and the semantic equivalence classes. The effect of strengthening an invariant is clear, and it shows the impact it may have on the resulting pre-processor.

We have shown here an example of a monolithic minimiser. The following example addresses the distributed case, when data comes from different sources.

7.3 Credit Score

This second example takes place from the point of view of *PrivaLoan*, a company specialised in making loans to individual customers. To decide whether or not loan is granted, *PrivaLoan* computes a credit score. This score depends on two elements: the credit history and the tax category of the individual. These two inputs are disclosed by two different parties: a central bank for the credit history and the tax office for the tax category.

One of the competitive advantage of *PrivaLoan* on the credit market is that it protects the privacy of the individuals asking for credit. This is achieved by applying the limited collection principle to personal data requested to get an

```
1  public class CreditApp {
2      //@ requires (0 <= incidents) && (incidents <= 3)
3      //@         && (1 <= tax)  && (tax <= 3);
4      public int compCreditScore(int incidents, int tax) {
5          int creditScore;
6          if (incidents == 0) {
7              if (tax == 3) { creditScore = 2; }
8              else { creditScore = 1; } }
9          else if (incidents == 1) { creditScore = 1; }
10         else { creditScore = 0; }
11         return creditScore; } }
```

Fig. 10: Program P_{cs} to compute a credit score.

Table 4: Semantics of P_{cs} (see Fig. 10).

variables		`incidents`			
	values	0	1	2	3
`tax`	1	1		0	
	2				
	3	2			

Note: There are three equivalence classes corresponding to the output values 0, 1, and 2.

offer. The rules to decide on the credit risk of an individual are implemented in the `compCreditScore` method of the program from Fig. 10. This method takes two arguments: `incidents`, which ranges from 0 to 3 as a grade for the credit-history of an individual, and `tax`, which ranges from 1 to 3 to indicate the tax category to which the individual belongs. For an incident level of 2 or 3, the credit score is 0, which means *PrivaLoan* does not want to make an offer. If the incident level is 0 or 1, the credit score is 1, except when the incident level is 0 and the tax category is 3, which leads to a credit score of 2.

As can be seen in Table 4, which maps the value of `creditScore` as function of `incidents` and `tax`, some combinations of inputs return the same score. Intuitively, there is no need for the central bank to make a distinction between incidents of grade 2 and 3 since the output is the same. Similarly, the tax office should not make a distinction between tax categories of 1 and 2. Indeed, for all the possible values of incidents, the credit score does not change between a tax category of 1 and one of 2 (though this score does not remain the same for all values of incidents).

The minimisers $M_{cs,0}$ and $M_{cs,1}$ in Fig. 11 and 12 are output by *DataMin* based on P_{cs} from Fig. 10. Both minimisers reflect the behavior described previously. As previously, we replaces the (weakest pre-)conditions by placeholders. Indeed, the representative assigned for the equivalence class $\{2, 3\}$ corresponds to the minimum value of this set (l.5 of the minimiser for the variable `incidents`, see Fig. 11). In all these cases, the value of the output of the program P_{cs} execution remains the same whichever the representative value chosen.

```
1 public class CreditAppMin_incidents {
2     public int minimise_compCreditScore(int incidents) {
3         int repr_incidents;
4         if (wp1) { repr_incidents = 1; }
5         else if (wp2) { repr_incidents = 2; }
6         else { repr_incidents = 0; }
7         return repr_incidents; } }
```

Fig. 11: First input (*incidents*) minimiser $M_{cs,0}$ for P_{cs} (see Fig. 10).

```
1 public class CreditAppMin_tax {
2     public int minimise_compCreditScore(int tax) {
3         int repr_tax;
4         if (wp1) { repr_tax = 1; }
5         else { repr_tax = 3; }
6         return repr_tax; } }
```

Fig. 12: Second input (*tax*) minimiser $M_{cs,1}$ for P_{cs} (see Fig. 10).

Similarly for the minimiser for the variable `tax` (see Fig. 12), the representative computed for the equivalence class $\{1, 2\}$ (l.4) can have any of these values without changing the value of the final output (1 has been fixed here).

In all the other cases, the representative must have the same value as the input variable (l.4 and l.6 of $M_{cs,0}$ and l.5 of $M_{cs,1}$). Indeed, without knowing a priori the value that will be disclosed by the other data source point, each has to assume it could be any value.

8 Related Work

Formal and rigorous approaches to privacy have long been advocated [30]. Many different dimensions in this field have been explored but the data minimisation principle seems not to have been precisely defined, as witnessed by [15], where it is stated that "a whole family of design principles are lumped under the term 'data minimisation"'. A related work is the notion of *minimal exposure* [2]. This consists in performing a preprocessing of information on the client's side to give only the data needed to benefit from a service. For instance, if the service to which an individual requests can be modeled as $a \vee b$, then if a is true it will be the only data disclosed: they reduce the *number* of inputs provided, but do not reason on the domain of the inputs.

The semantics of minimality and data minimisation is closely related to information flow, and we have used several ideas from that area [9, 14, 19, 28]. A semantic notion of dependency has been introduced in [22], where it is used to characterise the subset of variables in an expression containing all and only the variables that are semantically relevant for the evaluation of the expression. Our notion is related to this (viewing the variables of an expression as the input, and the value of the expression at the output), but much more fine-grained.

A notable difference in the formalisation of minimality compared to usual definitions of information flow is that minimality is a necessary and sufficient condition on information flow, whereas most security formulations of information flow are sufficient conditions only (e.g. noninterference: public inputs are sufficient to compute public outputs). An exception here is Chong's *necessary information release* [8] which provides both upper and lower bounds on information flows. For this reason many static analysis techniques that have been used for security properties (e.g. type systems) are not easy to use for minimality – an over-approximation to the information flows is not sound for minimality.

The necessary and sufficient conditions embodied in minimisers appeasr to be closely related to the notion of *completeness* in abstract interpretation [13], where a minimiser m plays the role of a *complete abstraction*. It is to be noted that some work in the domain of quantitative information flow aiming at automated discovery of leaks also rely on analysis of equivalence classes [5,21]. Compared to our proof-of-concept tool, the approach of [5] contains several ideas that could be beneficial to adopt in our tool.

The idea of equivalence partitioning by using symbolic execution has been used previously in the context of test case generation in [27]. Testing is also one of the uses of the symbolic execution as performed by the KeY theorem prover [1]. We use the KeY tool in our work to symbolically execute the given program, also exploiting the possibility of adding preconditions and invariants.

9 Conclusion

We have provided a formal definition of *data minimisation* in terms of strong dependency and derived concepts. We have also formalised the notion of a *data minimiser* as a pre-processor that may be composed with the data processor to filter the inputs so the latter receives representative inputs for different equivalence classes yielding the same output. We have considered the cases when the input comes from a unique source (monolithic collection), and when they come from different sources (distributed collection). We have also shown properties of both kinds of data minimisers depending on when they are used (before deployment, online, or offline).

Besides the theoretical results, we have also introduced a procedure, and a proof-of-concept implementation, to obtain data minimisers for a given program. The procedure is a chain consisting of a symbolic execution of the program and the use of a SAT solver to obtain a partition of the input space depending on the different possible outputs. The procedure generates representatives for each class which are outputs of the generated data minimiser to be fed as inputs into the data processor.

Given the nature of the problem our approach is semantics-based, so finding a distributed minimiser is undecidable in general. That said, as data has different sensitivity depending on the application, and the desirable privacy level, we could restrict the analysis to specific modules, making it tractable in some cases. The approach might also have practical utility if we aim at using it in an interactive

way, semi-automatically. One crucial issue in tackling realistic examples will be finding compositional methods. One possible compositional proof rule is the following: If m is a minimiser for p, n is a (distributed) minimiser for q and $\mathrm{range}(q) \subseteq \mathrm{range}(m)$ then n is a (distributed) minimiser for $p \circ q$.

Note the special case when p is minimal (and hence m is the identity) for which the condition $\mathrm{range}(q) \subseteq \mathrm{range}(m)$ becomes vacuous. It remains to be explored whether this has practical utility.

Our paper opens the way to other interesting research directions. On the theoretical side it would be worth exploring the connections to completeness in abstract interpretation as mentioned in the related work section. On the practical side we are considering the possibility of defining a type system to determine the cardinality of the needed inputs for computing a given output, and also the integration of a quantitative approach based on worth-based information theory to deal with the sensitivity of data.

References

1. Ahrendt, W., Baar, T., Beckert, B., Bubel, R., Giese, M., Hähnle, R., Menzel, W., Mostowski, W., Roth, A., Schlager, S., Schmitt, P.H.: The KeY tool. Software and System Modeling 4, 32–54 (2005)
2. Anciaux, N., Nguyen, B., Vazirgiannis, M.: Limiting data collection in application forms: A real-case application of a founding privacy principle. In: Privacy, Security and Trust (PST), 2012 Tenth Annual International Conference on. pp. 59–66 (July 2012)
3. Askarov, A., Sabelfeld, A.: Gradual release: Unifying declassification, encryption and key release policies. In: Proc. IEEE Symp. on Security and Privacy. pp. 207–221 (May 2007)
4. Askarov, A., Chong, S.: Learning is change in knowledge: Knowledge-based security for dynamic policies. In: Computer Security Foundations Symposium (CSF), 2012. pp. 308–322. IEEE (2012)
5. Backes, M., Kopf, B., Rybalchenko, A.: Automatic discovery and quantification of information leaks. In: Security and Privacy, 2009 30th IEEE Symposium on. pp. 141–153 (May 2009)
6. Balliu, M., Dam, M., Le Guernic, G.: Epistemic temporal logic for information flow security. ACM SIGPLAN Sixth Workshop on Programming Languages and Analysis for Security (PLAS 2011) (2011)
7. Bussard, L.: Encyclopedia of Cryptography and Security, chap. Secondary Use Regulations, pp. 1094–1095. Springer US (2011)
8. Chong, S.: Required information release. In: Computer Security Foundations Symposium (CSF), 2010 23rd IEEE. pp. 215–227 (July 2010)
9. Cohen, E.: Information transmission in computational systems. SIGOPS Oper. Syst. Rev. 11(5), 133–139 (Nov 1977), http://doi.acm.org/10.1145/1067625.806556
10. Dolzmann, A., Sturm, T.: Redlog: Computer algebra meets computer logic. SIGSAM Bull. 31(2), 2–9 (Jun 1997), http://doi.acm.org/10.1145/261320.261324
11. European Commission: Proposal for a Regulation of the European Parliament and of the Council on the protection of individuals with regard to the processing of personal data and on the free movement of such data (General Data Protection Regulation). Codecision legislative procedure for a regulation 2012/0011 (COD), European Commission, Brussels, Belgium (January 2012)
12. Flanagan, C., Qadeer, S.: Predicate abstraction for software verification. In: Proceedings of the 29th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages. pp. 191–202. POPL '02, ACM, New York, NY, USA (2002), http://doi.acm.org/10.1145/503272.503291
13. Giacobazzi, R., Ranzato, F.: Completeness in abstract interpretation: A domain perspective. In: AMAST. LNCS, vol. 1349. Springer (1997)
14. Giacobazzi, R., Mastroeni, I.: Abstract non-interference: Parameterizing non-interference by abstract interpretation. In: Proceedings of the 31st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages. pp. 186–197. POPL '04, ACM, New York, NY, USA (2004), http://doi.acm.org/10.1145/964001.964017
15. Gurses, S., Troncoso, C., Diaz, C.: Engineering privacy by design reloaded. Amsterdam Privacy Conference (2015)

16. Hentschel, M., Bubel, R., Hähnle, R.: Symbolic execution debugger (sed). In: Bonakdarpour, B., Smolka, S.A. (eds.) Proceedings of Runtime Verification 2014. pp. 255–262. LNCS, Springer (Sep 2014)
17. Hunt, L.S.: Abstract Interpretation of Functional Languages: From Theory to Practice. Ph.D. thesis, Department of Computing, Imperial College of Science, Technology and Medicine (1991)
18. King, J.C.: Symbolic execution and program testing. Communications of the ACM 19(7), 385–394 (Jul 1976), http://doi.acm.org/10.1145/360248.360252
19. Landauer, J., Redmond, T.: A lattice of information. In: Proc. IEEE Computer Security Foundations Workshop. pp. 65–70 (Jun 1993)
20. Leavens, G.T., Baker, A.L., Ruby, C.: Jml: a java modeling language. In: In Formal Underpinnings of Java Workshop (at OOPSLA'98 (1998)
21. Lowe, G.: Quantifying information flow. In: Proc. IEEE Computer Security Foundations Workshop. pp. 18–31 (Jun 2002)
22. Mastroeni, I., Zanardini, D.: Data dependencies and program slicing: From syntax to abstract semantics. In: Proceedings of the 2008 ACM SIGPLAN Symposium on Partial Evaluation and Semantics-based Program Manipulation. pp. 125–134. PEPM '08, ACM, New York, NY, USA (2008), http://doi.acm.org/10.1145/1328408.1328428
23. Mitchell, R., McKim, J., Meyer, B.: Design by Contract, by Example. Addison Wesley Longman Publishing Co., Inc., Redwood City, CA, USA (2002)
24. de Moura, L., Bjørner, N.: Z3: An efficient smt solver. In: Ramakrishnan, C., Rehof, J. (eds.) Tools and Algorithms for the Construction and Analysis of Systems, Lecture Notes in Computer Science, vol. 4963, pp. 337–340. Springer Berlin Heidelberg (2008), http://dx.doi.org/10.1007/978-3-540-78800-3_24
25. Organisation for Economic Co-operation and Development: Tech. rep., Organisation for Economic Co-operation and Development
26. Oystein, O.: Theory of equivalence relations. Duke Mathematical Journal 9, 573–627 (1942)
27. Richardson, D.J., Clarke, L.A.: A partition analysis method to increase program reliability. In: Proceedings of the 5th International Conference on Software Engineering. pp. 244–253. ICSE '81, IEEE Press, Piscataway, NJ, USA (1981), http://dl.acm.org/citation.cfm?id=800078.802537
28. Sabelfeld, A., Sands, D.: A per model of secure information flow in sequential programs. Higher-Order and Symbolic Computation 14(1), 59–91 (2001), http://dx.doi.org/10.1023/A%3A1011553200337
29. Smith, G.: On the foundations of quantitative information flow. In: de Alfaro, L. (ed.) Foundations of Software Science and Computational Structures, Lecture Notes in Computer Science, vol. 5504, pp. 288–302. Springer Berlin Heidelberg (2009), http://dx.doi.org/10.1007/978-3-642-00596-1_21
30. Tschantz, M.C., Wing, J.M.: FM 2009: Formal Methods: Second World Congress, Eindhoven, The Netherlands, November 2-6, 2009. Proceedings, chap. Formal Methods for Privacy, pp. 1–15. Springer Berlin Heidelberg, Berlin, Heidelberg (2009), http://dx.doi.org/10.1007/978-3-642-05089-3_1
31. US Secretary's Advisory Committee on Automated Personal Data Systems: Records, Computers and the Rights of Citizens. Report DHEW NO. (OS)73-94, US Secretary's Advisory Committee on Automated Personal Data Systems, Brussels, Belgium (July 1973), chapter IV: Recommended Safeguards for Administrative Personal Data Systems

Appendix A: Proof of Theorem 3

Let $\mathrm{ER}(D)$ denote the equivalence relations over set D.

We view equivalence relations as sets of pairs, writing either aRb or $(a,b) \in R$ when relation R relates a to b. Note that $\mathrm{ER}(D)$ forms a complete lattice $\langle \mathrm{ER}(D), \bot, \top, \leq, \vee, \wedge \rangle$ where $\bot = \{(a,a) \mid a \in D\}$, $\top = D \times D$, $\leq$ is subset inclusion $\subseteq$, $\vee$ is transitive closure of set union $\cup^*$, and $\wedge$ is set intersection $\cap$.

Definition 9. *Given $R \in \mathrm{ER}(D)$ and $S \in \mathrm{ER}(E)$ let $R \boxtimes S \in \mathrm{ER}(D \times E)$ denote the relation $(d,e)R \boxtimes S(d',e') \Leftrightarrow dRd' \wedge eSe'$.*

The lattice of equivalence relations $\mathrm{ER}(D)$ forms a complete lattice (standard) as described above.

Proposition 5. $\left\{ \boxtimes_{i \in \{0,\ldots,n\}} R_i \mid R_i \in \mathrm{ER}(D_i) \right\}$ *is a complete sub-lattice of* $\mathrm{ER}\left(\prod_{i \in \{0,\ldots,n\}} D_i \right)$.

Proof. Note that in $\mathrm{ER}(D)$, $\bot$ is the identity on D, written Id_D and $\top$ is the universal relation All_D.

It is easy to see that $\boxtimes_{i \in \{0,\ldots,n\}} \mathrm{Id}_{D_i} = \mathrm{Id}_{\left(\boxtimes_{i \in \{0,\ldots,n\}} D_i \right)}$ and $\boxtimes_{i \in \{0,\ldots,n\}} \mathrm{All}_{D_i} = \mathrm{All}_{\left(\boxtimes_{i \in \{0,\ldots,n\}} D_i \right)}$.

Similarly, it is straightforward to show that

$\boxtimes_{i \in \{0,\ldots,n\}} R_i \cup \boxtimes_{i \in \{0,\ldots,n\}} S_i = \boxtimes_{i \in \{0,\ldots,n\}} (R_i \cup S_i)$ and similarly for $\cap$ (by observing that union, intersection, and transitive closure all act pointwise on t-uples).

PROOF OF THE THEOREM

We need to prove that *for any* $dp \in \left(\prod_{i \in \{0,\ldots,n\}} I_i \right) \to O$, $\exists dm = \otimes_{i \in \{0,\ldots,n\}} m_i$ *such that dm is a best distributed minimiser for dp.*

Proof. We first build a candidate *dm*, and by assuming that it is *not* a best distributed minimiser, derive a contradiction.

Let $E = \left\{ \boxtimes_{i \in \{0,\ldots,n\}} R_i \mid R_i \in \mathrm{ER}(I_i) \right\}$. Note that E is a complete sub-lattice of $\mathrm{ER}\left(\prod_{i \in \{0,\ldots,n\}} I_i \right)$ (by Proposition 5).

Define $K = \cup^* \{ R \in E \mid R \subseteq \ker dp \}$.

By the sub-lattice property of E above, $K \in E$, ie. $K = \prod_{i \in \{0,\ldots,n\}} K_i$ for some $K_i \in \mathrm{ER}(I_i)$ (for $i \in \{0,\ldots,n\}$) and $K \subseteq \ker dp$ (since $\ker dp$ is an upper bound, it is also an upper bound for the least upper bound).

Now we define a candidate $dm = \bigotimes_{i \in \{0,\ldots,n\}} m_i$. To do so, for each equivalence class C of each K_i (for $i \in \{0,\ldots,n\}$), choose a representative element denoted $\mathrm{repr}(C)$.

Now define $m_i(v) \triangleq \mathrm{repr}\left([v]_{K_i}\right)$, ie. m_i maps a value into the representative chosen for that value's equivalence class under K_i.

First we show that dm is a minimiser for P.

$$\begin{aligned} dp(dm(v)) &= dp\left(m_0(v_0), \ldots, m_n(v_n)\right) \\ &= dp\left(\mathrm{repr}\left([v_0]_{K_0}\right), \ldots, \mathrm{repr}\left([v_n]_{K_n}\right)\right) \end{aligned}$$

Note that $\left(\mathrm{repr}\left([v_i]_{K_i}\right), v_i\right) \in K_i$ (for $i \in \{0,\ldots,n\}$) then $\left(\mathrm{repr}\left([v_0]_{K_0}\right), \ldots, \mathrm{repr}\left([v_n]_{K_n}\right)\right)$ is related to v by K and hence by $\ker dp$.

Thus $dp\left(\mathrm{repr}\left([v_0]_{K_0}\right), \ldots, \mathrm{repr}\left([v_n]_{K_n}\right)\right) = dp(v)$.

It remains to show that dm is idempotent. This follows easily from the fact that the m_i are idempotent since $\mathrm{repr}(C) \in C$ for any equivalence class of m_i.

Finally, assume towards a contradiction that dm thus defined is *not* a best distributed minimiser for dp.

Let $J_i = \mathrm{range}\left(m_i\right)$. Then $\mathrm{range}\left(m_i\right) = \prod_{i \in \{0,\ldots,n\}} J_i$. Then by the assumption dp is *not* best distributed-minimised for $\prod_{i \in \{0,\ldots,n\}} J_i$, $\exists a \in \{0,\ldots,n\}$ and distinct values $u, v \in J_a$ such that $\forall \overline{u}, \overline{v}$ which differ only at position a, where $\overline{u}_a = u, \overline{v}_a = v$ we have $dp\left(\overline{u}\right) = dp\left(\overline{v}\right)$.

Define $R_a = \{(u,v),(v,u)\} \cup^* \mathrm{Id}_{J_a}$. Define $R = \boxtimes_{i \in \{0,\ldots,n\}} R_i$ where $R_i = \mathrm{Id}_{J_i}$ for $i \neq a$.

Thus $(\overline{u}, \overline{v}) \in R \Rightarrow dp\left(\overline{u}\right) = dp\left(\overline{v}\right)$ and hence $R \subseteq \ker dp\,|_{\mathrm{range}(dm)}$, and hence (*) $R \subseteq K$.

By construction of m_a, since $u \neq v$, then u and v must not be related by K_a which contradicts (*).